

Realizing Grid Technology: Roadblocks and Solutions

Dennis Gannon, Dept. of Computer Science, Indiana University

Grids are designed to provide a way for large scientific collaborations to increase productivity by sharing resources. This idea has been both a success and a failure. Where it has succeeded is in areas of high energy physics, where the need for a large group to collaborate is absolutely essential: they must share access to a large instrument that is the focus of all of their work. Where it has failed is in the case need to collaborate is perceived to be less demanding or in those cases where security and social barriers have been so large that installing Grid software that crosses institutional boundaries is nearly impossible. Without strong pressure from the collaborating scientists to redirect the institutional inertia, the Grid does not happen.

Unfortunately, this problem is becoming worse and not better. Cyber terrorism has created a mindset that forces institutions to retreat behind larger and larger firewalls. Securing computer accounts for collaborators outside ones institution is becoming increasingly difficult to do. As a result, our current approach to building Grids based on sharing access to computer systems simply does not scale.

However, there are several areas where distributed systems technology is growing or is well established and may point the way to building stronger Grid systems that can help a greater number of scientists. For example, the use of web technologies and the Access Grid in scientific collaboration continue to grow at a steady rate. Every scientist at all but the most restricted laboratories has a way to publish scientific papers, software and data using a website. Widely used web applications like “blogs” provide a journaling capability that is not far from the goals of the DOE scientific notebooks. Furthermore, peer-to-peer technologies for data sharing are in widespread use and their application goes well beyond the sharing of copyrighted media. Finally powerful search technology has revolutionized scientific information sharing.

The reason that this technology works where “traditional” Grid solutions are stumbling is very simple. These technologies are based on providing services and not servers. Services are software systems that provide users capabilities without requiring that the clients have direct access to the back-end hardware. A service can be installed within a firewall and, if properly vetted by institutional security authorities, the service can be exposed to external users through a gateway portal. For example, a scientist can provide a simulation tool that may be accessed via a user interface in a portal by any member of an authenticated and authorized virtual organization. The simulation can run using institutional resources managed under a single account that can be carefully monitored and controlled. Such a service may also be accessed by other services as part of a distributed Grid workflow provided a valid identity has been delegated to the client services.

These facts have not been lost on the Grid research community. A web service model is the foundation of all new Grid software technologies such as Globus and OGSA. The current DOE portal work is all web service based. However, there are several important missing pieces.

First, for scientists to make use of this technology it must be easy to wrap an application as a secure service. In some cases the application is a simulation but in other cases it may be an archive of experimental results that are to be made available to authorized users. What tools are available for a scientist to advertise the content of an archive or interface to an application so that it can be discovered by only authorized users? Can we build a standard “science service google” that can be used to crawl service registries and parse and catalog scientific service metadata? How do we build easy-to-use tools to automatically “wrap” applications and archives as services?

Second, how does a “client” scientist write programs or workflows that use scientific web services as part of larger computation? Fortunately, web service programming tools are becoming more widely available. Python and Perl interfaces to secure web services are available and these also adhere to Grid standards such as GSI. Desktop workflow tools like Triana and Taverna allow web services to be incorporated into simple desktop computations. Even Matlab now has an API to call web services.

Another limitation of current web service technology is its poor performance profile. Can we work with industry to provide a way for large-scale scientific data to be exchanged using as WS framework but without the overhead of SOAP and XML encodings. GridFTP is one solution, but FTP is not the most appropriate protocol for many applications (such as streaming data).

To achieve scalability we need to be able to allow a more peer-to-peer Grid model evolve that will allow Grid scientific technology to be easily deployed by an individual on a desktop computer. Because web service technology is lightweight and very portable, this is a reachable goal. We need to provide three additional components:

- A lightweight P2P authorization system that make it easy for a programmer to provide authorization capabilities to collections of remote users.
- A well-established metadata publishing/discovery system so that an individual can describe and advertise services to authorized communities.
- A better data transport protocol that works within a WS framework but is more appropriate for scientific data.

Many others have described similar ideas for P2P and “lightweight” Grid tools. It may be time for DOE to take a coordinated approach to providing them.